arXiv:1703.00542v1 [math.ST] 1 Mar 2017

A Note on the Approximate Admissibility of Regularized Estimators in the Gaussian Sequence Model

Xi Chen[1], **Adityanand Guntuboyina**[2,*] **and Yuchen Zhang**[3]

[1]*Stern School of Business*
New York University
New York, New York, 10012
e-mail: xichen@nyu.edu

[2]*Department of Statistics*
University of California, Berkeley
Berkeley, CA, 94720
e-mail: aditya@stat.berkeley.edu

[3]*Computer Science Department*
Stanford University
Stanford, CA, 94305
e-mail: zhangyuc@cs.stanford.edu

Abstract: We study the problem of estimating an unknown vector θ from an observation X drawn according to the normal distribution with mean θ and identity covariance matrix under the knowledge that θ belongs to a known closed convex set Θ. In this general setting, Chatterjee (2014) proved that the natural constrained least squares estimator is "approximately admissible" for every Θ. We extend this result by proving that the same property holds for all convex penalized estimators as well. Moreover, we simplify and shorten the original proof considerably. We also provide explicit upper and lower bounds for the universal constant underlying the notion of approximate admissibility.

Keywords and phrases: Admissibility, Bayes risk, Gaussian sequence model, Least squares estimator, Minimaxity.

1. Introduction

The Gaussian sequence model is a commonly used model for theoretical investigations in nonparametric and high dimensional statistical problems. Here one models the data vector $X \in \mathbb{R}^n$ as an observation having the normal distribution with unknown mean $\theta \in \mathbb{R}^n$ and identity covariance matrix i.e., $X \sim N(\theta, I_n)$. Often one assumes some structure on the unknown mean θ in the form of a convex constraint. Specifically, it is common to assume that $\theta \in \Theta$ for some closed convex subset Θ of $\mathbb{R}^n$. A natural estimator for θ under the constraint $\theta \in \Theta$ is the least squares estimator (LSE) defined as

$$\widehat{\theta}(X;\Theta) := \operatorname*{argmin}_{\alpha \in \Theta} \frac{1}{2}\|X - \alpha\|_2^2 \tag{1.1}$$

*Supported by NSF Grant DMS-1309356

where $\|\cdot\|_2$ denotes the usual Euclidean norm on $\mathbb{R}^n$. It is easy to see that many common estimators in nonparametric and high-dimensional statistics such as shape constrained estimators (see, for example, Groeneboom and Jongbloed (2014)) and those based on constrained LASSO (see, for example, Bühlmann and van De Geer (2011)) are special cases of the LSE (1.1) for various choices of Θ.

In this abstract setting, Chatterjee (2014) asked the following question: Does the estimator $\widehat{\theta}(X;\Theta)$ satisfy a general optimality property that holds for *every* closed convex set Θ? This is a non-trivial question; obvious guesses for the optimality property might be admissibility and minimaxity but the LSE does not satisfy either of these for every Θ. Indeed, $\widehat{\theta}(X;\Theta)$ is not minimax (even up to multiplicative factors that do not depend on the dimension n) when $\Theta := \{\alpha \in \mathbb{R}^n : \sum_{i<n} \alpha_i^2 + n^{-1/2}\alpha_n^2 \le 1\}$ as noted by Zhang (2013) (a more elaborate counterexample for minimaxity is given in Chatterjee (2014)). Also, $\widehat{\theta}(X;\Theta)$ is not admissible when $\Theta = \mathbb{R}^n$ where the James-Stein estimator dominates $\widehat{\theta}(X) = X$ (see, for example, Lehmann and Casella (1998)).

Chatterjee (2014) answered the general optimality question of the constrained LSE in the affirmative by proving that $\widehat{\theta}(X;\Theta)$ is *approximately admissible* over Θ for every Θ. The precise statement of Chatterjee's theorem is described below. Let us say that, for a constant $C > 0$, an estimator $\mathfrak{d}(X)$ is *C-admissible* over Θ if for every other estimator $\widetilde{\mathfrak{d}}(X)$, there exists $\theta \in \Theta$ such that

$$C\mathbb{E}_\theta\|\mathfrak{d}(X) - \theta\|_2^2 \le \mathbb{E}_\theta\|\widetilde{\mathfrak{d}}(X) - \theta\|_2^2. \tag{1.2}$$

In words, the above definition means that for every estimator $\widetilde{\mathfrak{d}}(X)$, there exists a point $\theta \in \Theta$ at which the estimator $\mathfrak{d}(X)$ performs as well as the estimator $\widetilde{\mathfrak{d}}(X)$ up to the multiplicative factor C. Note that the point at which $\mathfrak{d}(X)$ performs better than $\widetilde{\mathfrak{d}}(X)$ would depend on the estimator $\widetilde{\mathfrak{d}}(X)$ as well as on the constraint set Θ. Essentially an estimator $\mathfrak{d}(X)$ being C-admissible over Θ means that it is impossible for any estimator to dominate $\mathfrak{d}(X)$ uniformly over Θ by more than the multiplicative factor C.

Chatterjee (2014) proved that there exists a *universal* constant $0 < C \le 1$ such that for every $n \ge 1$ and closed convex subset $\Theta \subseteq \mathbb{R}^n$, the LSE $\widehat{\theta}(X;\Theta)$ is C-admissible for Θ.

Theorem 1.1. [Chatterjee (2014)] There exists a universal constant $0 < C \le 1$ (independent of n and Θ) such that for every $n \ge 1$ and closed convex subset $\Theta \subseteq \mathbb{R}^n$, the least squares estimator $\widehat{\theta}(X;\Theta)$ is C-admissible over Θ.

Remarkable features of the above theorem are that it is true for every Θ and that the constant C does not depend on n or Θ. We would like to mention here that Theorem 1.1 is a rather difficult result (in Chatterjee's own words, "from a purely mathematical point of view, this is the deepest result of this paper") and the original proof in Chatterjee (2014) is quite complex.

Our paper has the following twin goals: (a) we extend Theorem 1.1 (which only involves constrained estimators) to penalized estimators, which are more commonly used in practice, and (b) we simplify considerably the proof of Theorem 1.1 given in Chatterjee (2014) and our proof is also much more intuitive.

To describe our main result, let us first introduce penalized estimators. Given a closed convex set $\Theta \subseteq \mathbb{R}^n$ and a real-valued convex function f on Θ, let

$$\widehat{\theta}(X;\Theta,f) := \operatorname*{argmin}_{\alpha\in\Theta}\left(\frac{1}{2}\|X-\alpha\|^2 + f(\alpha)\right). \tag{1.3}$$

Strictly speaking $\widehat{\theta}(X;\Theta,f)$ is a least squares estimator that is both constrained and penalized. We can of course write it as a pure penalized estimator with the penalty function $\tilde{f}(x) = f(x)+\mathbb{I}_\Theta(x)$, where $\mathbb{I}_\Theta(x)$ is the indicator function that takes the value 0 when $x \in \Theta$ and $+\infty$ otherwise. We choose to separate the constraint and penalty as it is more natural for many statistical applications. In doing so, note that we have required that f is real-valued (i.e., f does not take the value $+\infty$) on Θ.

For $\Theta = \mathbb{R}^n$ in (1.3), we obtain penalized estimators for which the LASSO is the most common example. For $f \equiv 0$, we get back the constrained LSEs of (1.1). There are examples where one uses both a non-trivial constraint set Θ and a non-trivial penalty function $f(\cdot)$; for example, in isotonic regression, it is common to use

$$\Theta := \{\alpha \in \mathbb{R}^n : \alpha_1 \le \cdots \le \alpha_n\} \quad \text{and} \quad f(\alpha) := \lambda\,(\alpha_n - \alpha_1)$$

for some $\lambda \ge 0$. This estimator fits non-decreasing sequences to the data while constraining the range of the estimator so as to prevent the spiking effect that the usual isotonic LSE suffers from; see, for example, Woodroofe and Sun (1993).

Because of the presence of the penalty function f, it is clear that the class of estimators given by $\widehat{\theta}(X;\Theta,f)$ is larger compared to the class given by the LSEs in (1.1). The main result of our paper is the following.

Theorem 1.2. There exists a universal constant $0 < C \le 1$ (independent of n, Θ and f) such that for every $n \ge 1$, closed convex set $\Theta \subseteq \mathbb{R}^n$ and real-valued convex function f on Θ, the estimator $\widehat{\theta}(X;\Theta,f)$ is C-admissible over Θ.

The above theorem generalizes Theorem 1.1 by showing that all estimators $\widehat{\theta}(X;\Theta,f)$ have the C-admissibility property over Θ for a universal constant C. In words, this means that given any estimator $\mathfrak{d}(X)$, there exists a point $\theta \in \Theta$ at which the estimator $\widehat{\theta}(X;\Theta,f)$ performs as well as the estimator $\mathfrak{d}(X)$ up to the multiplicative factor C. This point $\theta \in \Theta$ would depend on the estimator $\mathfrak{d}(X)$ as well as on the constraint set Θ and the penalty function f.

It should be noted here that C-admissibility (even 1-admissibility) does not by itself guarantee than an estimator is good in a reasonable sense. This is because unnatural estimators that return a fixed vector in the parameter space are admissible. However, C-admissibility could serve as a minimum requirement for an estimator to be reasonable. Our main point here is that the estimators $\widehat{\theta}(X;\Theta,f)$ are very natural and commonly used in many applications. Our result shows that all these natural estimators (as Θ and f varies) satisfy C-admissibility for a universal constant C.

1.1. Connections to the normalized minimax risk

There is a restatement of Theorem 1.2 that is illuminating and gives a minimax flavor to Theorem 1.2. Given Θ and f, let us define the normalized minimax risk over Θ by

$$R_{\text{nor}}(\Theta; f) := \inf_{\mathfrak{d}} \sup_{\theta \in \Theta} \frac{\mathbb{E}_\theta \|\mathfrak{d}(X) - \theta\|_2^2}{\mathbb{E}_\theta \|\widehat{\theta}(X; \Theta, f) - \theta\|_2^2}. \tag{1.4}$$

where the infimum is over all estimators $\mathfrak{d}(X)$. Here we use the conventions $\frac{0}{0} = 1$ and $\frac{a}{0} = +\infty$ for $a > 0$. Note that $R_{\text{nor}}(\Theta; f)$ is defined just like the usual minimax risk over the parameter space Θ except that the risk of every estimator $\mathfrak{d}(X)$ is rescaled (normalized) by the risk of $\widehat{\theta}(X; \Theta, f)$. This therefore a reasonable measure of comparison of arbitrary estimators $\mathfrak{d}(X)$ to our estimator $\widehat{\theta}(X; \Theta, f)$.

It is clear that $R_{\text{nor}}(\Theta; f) \leq 1$ as can be seen by bounding the infimum in (1.4) by the term corresponding to $\mathfrak{d}(X) = \widehat{\theta}(X; \Theta, f)$. A small value for $R_{\text{nor}}(\Theta; f)$ means that there exists an estimator $\mathfrak{d}(X)$ and some point $\theta \in \Theta$ at which the risk of $\mathfrak{d}(X)$ is smaller, by a large factor, than the risk of $\widehat{\theta}(X; \Theta, f)$.

Let us now define a universal constant C^* by taking the *worst possible value* of $R_{\text{nor}}(\Theta; f)$ over all possible values of the dimension n, convex constraint set Θ and convex penalty function f. Specifically, let

$$C^* := \inf_{n \geq 1} \inf_{\Theta \in \mathcal{C}_n} \inf_{f \in \mathcal{F}(\Theta)} R_{\text{nor}}(\Theta; f) \tag{1.5}$$

where $\mathcal{C}_n$ denotes the class of all closed convex subsets of $\mathbb{R}^n$ and $\mathcal{F}(\Theta)$ denotes the class of all real-valued convex functions on Θ. Note first that C^* is a universal constant and, a priori, it is not clear if C^* is zero or strictly positive.

It is now straightforward to verify that Theorem 1.2 is equivalent to the statement that C^* *is strictly positive*. Another contribution of our paper is to provide explicit lower and upper bounds for C^*.

Theorem 1.3. The universal constant C^* satisfies

$$6.05 \times 10^{-6} \leq C^* \leq \frac{1}{2}. \tag{1.6}$$

The lower bound of 6.05×10^{-6} for C^* comes from our argument for the proof of Theorem 1.2. It must be noted here that Chatterjee (2014) does not provide any explicit values for C in his C-admissibility result. Even if the constant C were tracked down in the proof of Chatterjee (2014), it appears that it will be smaller than 6.05×10^{-6} by several orders of magnitude. The improvement of the lower bound also shows the advantage of using our new arguments in the proof of admissibility.

The upper bound of $1/2$ for C^* is a consequence of an explicit construction of Θ and f such that $\widehat{\theta}(X; \Theta, f)$ is uniformly dominated over Θ by a factor of 2 by another estimator. We believe that this example is non-trivial. Please see Section 4 for the proof of Theorem 1.3.

The determination of the exact value of the constant C^* is likely to be a very challenging problem, which is left to be a future work.

1.2. Proof sketch

As a summary, the contributions of the paper include: (1) a novel and intuitive proof of a generalization of a result of Chatterjee (2014) on C-admissibility in Theorem 1.2; (2) explicit bounds for the worst possible value of the normalized minimax risk C^* in Theorem 1.3. In this subsection, we provide an outline of our proofs of these results.

Admissibility results are almost always proved via Bayesian arguments involving priors. Analogous to the notion of C-admissibility, we can define a notion of C-Bayes as follows. For $C > 0$ and a proper prior w over Θ, we say that an estimator $\mathfrak{d}(X)$ is C-Bayes with respect to w if

$$C \int_\Theta \mathbb{E}_\theta \|\mathfrak{d}(X) - \theta\|_2^2 w(d\theta) \le R_{\text{Bayes}}(w) := \inf_{\tilde{\mathfrak{d}}} \int_\Theta \mathbb{E}_\theta \|\tilde{\mathfrak{d}}(X) - \theta\|_2^2 w(d\theta) \quad (1.7)$$

where $R_{\text{Bayes}}(w)$ is the Bayes risk with respect to w and the infimum in the definition of $R_{\text{Bayes}}(w)$ above is over all estimators $\tilde{\mathfrak{d}}$.

It is now trivial to see that an estimator $\mathfrak{d}(X)$ is C-admissible over Θ if it is C-Bayes for some proper prior w supported on Θ. As a result, in order to prove that $\widehat{\theta}(X; \Theta, f)$ is C-admissible over Θ, it is sufficient to construct a proper prior w on Θ such that $\widehat{\theta}(X; \Theta, f)$ is C-Bayes with respect to w. We construct such a prior w by modifying the construction of Chatterjee (2014) (which only applied to the LSE) appropriately. See the proof of Theorem 1.2 for the description of w.

For our prior w, in order to prove that $\widehat{\theta}(X; \Theta, f)$ is C-Bayes, we need to

1. bound $\int_\Theta \mathbb{E}_\theta \|\widehat{\theta}(X; \Theta, f) - \theta\|_2^2 w(d\theta)$ from above, and
2. bound $R_{\text{Bayes}}(w)$ from below

and make sure that the two bounds differ only by the multiplicative factor C. For the first step above, we need to study the risk properties of the estimator $\widehat{\theta}(X, \Theta, f)$. This is done in Section 2. For the second step, we apply a recent general Bayes risk lower bound from Chen et al. (2016). The application of this risk bound shortens the proof considerably. In contrast, Chatterjee (2014) used a bare hands for lower bounding the Bayes risk via "a sequence of relatively complicated technical steps involving concentration inequalities and second moment lower bounds".

For proving Theorem 1.3, we first observe that our proof of Theorem 1.2 also yields the lower bound of 6.05×10^{-6} for the constant C^*. For proving that $C^* \le 1/2$, we explicitly construct a convex set Θ over which the normalized minimax risk is arbitrarily close to $1/2$ (see Section 4 for details).

The rest of this paper is structured as follows. In Section 2, we describe some results on the risk of the estimator $\widehat{\theta}(X; \Theta, f)$. These can be seen as an extension of the results of Chatterjee (2014) for penalized estimators. We will also discuss

the connection of our risk bounds to a recent work by van de Geer and Wainwright (2015). Section 3 contains the proof of Theorem 1.2 while Section 4 contains the proof of Theorem 1.3. Section 5 contains the proofs for the risk results of penalized estimators from Section 2.

2. Risk Behavior of $\widehat{\theta}(X;\Theta,f)$

Throughout this section, we fix a closed convex set Θ in $\mathbb{R}^n$ and a real-valued convex function f on Θ. The data vector X will be generated according to the normal distribution with mean θ and identity covariance. We study the risk of the estimator $\widehat{\theta}(X;\Theta,f)$. The main risk result is Theorem 2.1 below which will be used in the proof of Theorem 1.2 to bound the quantity

$$\int_{\Theta} \mathbb{E}_\theta \|\widehat{\theta}(X;\Theta,f)-\theta\|_2^2 w(d\theta)$$

for a suitable prior w.

The basic fact about the estimator $\widehat{\theta}(X;\Theta,f)$ (proved in Theorem 2.1 below) is that the loss $\|\widehat{\theta}(X;\Theta,f)-\theta\|_2$ is concentrated around a deterministic quantity t_θ which depends on θ, the constraint set Θ and the regularizer f. The quantity t_θ is defined as the maximizer of the function $G_\theta : [0,\infty) \to \mathbb{R}$ over $[0,\infty)$ where

$$G_\theta(t) := m_\theta(t) - \frac{t^2}{2} \qquad \text{with } m_\theta(t) := \mathbb{E}_\theta \left(\sup_{\alpha\in\Theta:\|\alpha-\theta\|_2\le t} \{\langle X-\theta, \alpha-\theta\rangle - f(\alpha)\} \right). \tag{2.1}$$

The quantity $m_\theta(t)$ can be viewed as an extension of notion of (localized) Gaussian width with the penalty function $f(\alpha)$ (note that $X-\theta$ is a standard Gaussian random variable). Indeed, when $f \equiv 0$, $m_\theta(t)$ is the Gaussian width of the set $\{\alpha-\theta : \alpha\in\Theta, \|\alpha-\theta\|_2 \le t\}$. The existence of t_θ as a unique maximizer of $G_\theta(t)$ over $t\in[0,\infty)$ is proved in Lemma 2.2 (see the end of this section). We also note that t_θ depends on the choice of the penalty f.

Theorem 2.1. Fix $\theta\in\Theta$ and consider the estimator $\widehat{\theta}(X;\Theta,f)$ constructed from X generated according to the model $X\sim N(\theta, I_n)$. Then

$$\mathbb{P}\left\{ \|\widehat{\theta}(X;\Theta,f)-\theta\|_2 \ge t_\theta + \delta \right\} \le 2\exp\left(-\frac{\delta^4}{32(t_\theta+\delta)^2} \right) \tag{2.2}$$

for every $\delta\ge 0$ and

$$\mathbb{E}_\theta \|\widehat{\theta}(X;\Theta,f)-\theta\|_2^2 \le t_\theta^2 + \left(2\sqrt{84}\right) t_\theta \min(\sqrt{t_\theta},1) + 84\min(t_\theta,1). \tag{2.3}$$

When $f\equiv 0$ i.e., when the estimator $\widehat{\theta}(X;\Theta,f)$ becomes the LSE over Θ, then the above result has been proved by Chatterjee (2014, Theorem 1.1). Therefore, Theorem 2.1 can be seen as an extension of Chatterjee (2014, Theorem 1.1) to penalized estimators. Muro and van de Geer (2015) also studied concentration

for penalized estimators; however their result (see Muro and van de Geer (2015, Theorem 1)) proves concentration for (in our notation) the quantity

$$\tau(\widehat{\theta}(X;\Theta,f)) := \sqrt{\|\widehat{\theta}(X;\Theta,f)-\theta\|_2^2 + 2f(\widehat{\theta}(X;\Theta,f))}.$$

More recently, van de Geer and Wainwright (2015) studied concentration of the loss of empirical risk minimization estimators in a very general setting. Two of their results are relevant to Theorem 2.1. In van de Geer and Wainwright (2015, Theorem 2.1), it is proved that $\|\widehat{\theta}(X;\Theta,f)-\theta\|_2$ concentrates around its expectation $\mathbb{E}_\theta\|\widehat{\theta}(X;\Theta,f)-\theta\|_2$ at a rate that is faster than that given by Theorem 2.1. However to prove our admissibility result, we require concentration of $\|\widehat{\theta}(X;\Theta,f)-\theta\|_2$ around t_θ and not around $\mathbb{E}_\theta\|\widehat{\theta}(X;\Theta,f)-\theta\|_2$. The relation between t_θ and $\mathbb{E}_\theta\|\widehat{\theta}(X;\Theta,f)-\theta\|_2$ is not completely clear. Another result from van de Geer and Wainwright (2015) that is relevant to us is their Theorem 4.1. However it also gives concentration for the quantity $\tau(\widehat{\theta}(X;\Theta,f))$ while we require concentration for $\|\widehat{\theta}(X;\Theta,f)-\theta\|_2$. It is also worthwhile to note that van de Geer and Wainwright (2015) also studied concentration in models more general than Gaussian sequence models.

In addition to Theorem 2.1, we shall require some additional facts about t_θ and the function m_θ. These are summarized in the following result which also includes a statement on the existence and uniqueness of t_θ. For the case of the LSE (i.e., when $f \equiv 0$), the facts stated in the lemma below are observed in Chatterjee (2014) and most of the results in the following lemma are straightforward extensions of the corresponding facts in Chatterjee (2014).

Lemma 2.2. Recall the functions $G_\theta(\cdot)$ and $m_\theta(\cdot)$ from (2.1).

1. For every $\theta \in \Theta$, the function m_θ is non-decreasing and concave.
2. For every $\theta \in \Theta$, the function $G_\theta(\cdot)$ has a unique maximizer t_θ on $[0,\infty)$.
3. For every $\theta \in \Theta$ and $t \geq 0$, we have
$$m_\theta(t) \leq m_\theta(t_\theta) + t_\theta(t-t_\theta). \tag{2.4}$$
4. For every $\theta \in \Theta$ and $t \geq 0$, we have
$$G_\theta(t) - G_\theta(t_\theta) \leq \frac{-(t-t_\theta)^2}{2}. \tag{2.5}$$
5. The risk function $\theta \mapsto \mathbb{E}_\theta\|\widehat{\theta}(X;\Theta,f)-\theta\|_2^2$ is smooth in the following sense: for every $\theta_1,\theta_2 \in \Theta$, we have
$$\mathbb{E}_{\theta_1}\|\widehat{\theta}(X;\Theta,f)-\theta_1\|_2^2 \leq 2\mathbb{E}_{\theta_2}\|\widehat{\theta}(X;\Theta,f)-\theta_2\|_2^2 + 8\|\theta_1-\theta_2\|^2. \tag{2.6}$$
6. The function $\theta \mapsto t_\theta$ is smooth in the following sense: for every $\theta_1,\theta_2 \in \Theta$, we have
$$\left(t_{\theta_1} - \sqrt{\|\theta_1-\theta_2\|_2^2 + 4t_{\theta_1}\|\theta_1-\theta_2\|_2}\right)_+ \leq t_{\theta_2} \leq t_{\theta_1} + \sqrt{\|\theta_1-\theta_2\|_2^2 + 4t_{\theta_1}\|\theta_1-\theta_2\|_2}. \tag{2.7}$$

Also for every $\theta_1, \theta_2 \in \Theta$ and $\rho \geq 0$, we have

$$\left(1 - \sqrt{\rho^2 + 4\rho}\right)_+ t_{\theta_1} \leq t_{\theta_2} \leq \left(1 + \sqrt{\rho^2 + 4\rho}\right) t_{\theta_1} \tag{2.8}$$

provided $\|\theta_1 - \theta_2\|_2 \leq \rho t_{\theta_1}$.

With Theorem 2.1 and Lemma 2.2 in place, we first provide the proof of our main result (Theorem 1.2) in the next section. The proofs for Theorem 2.1 and Lemma 2.2 will be relegated to the end of the paper in Section 5.

3. Proof of Theorem 1.2

We follow the program outlined in the introduction. For proving that $\widehat{\theta}(X; \Theta, f)$ is C-admissible for a constant C, it is enough to demonstrate the existence of a prior w on Θ such that $\widehat{\theta}(X; \Theta, f)$ is C-Bayes with respect to w. As described in the introduction, a key step for proving that $\widehat{\theta}(X; \Theta, f)$ is C-Bayes involves bounding from below the Bayes risk $R_{\text{Bayes}}(w)$ with respect to w. For this purpose, we shall use the following result from Chen et al. (2016, Corollary 4.4). This result states that the following inequality holds for every prior w on Θ:

$$R_{\text{Bayes}}(w) \geq \frac{1}{2} \sup \left\{ t > 0 : \sup_{a \in \Theta} w\{\theta \in \Theta : \|\theta - a\|_2^2 \leq t\} < \frac{1}{4(1+I)} \right\} \tag{3.1}$$

where I is any nonnegative number satisfying

$$I \geq \inf_Q \int_\Theta \chi^2(P_\theta \| Q) dw(\theta). \tag{3.2}$$

Here P_θ denotes the n-dimensional normal distribution with mean θ and identity covariance and the infimum in (3.2) is over all probability measures Q on $\mathbb{R}^n$. Also $\chi^2(P\|Q)$ denotes the chi-square divergence defined as $\int (p^2/q) d\mu - 1$ where p and q are densities of P and Q respectively with respect to a common dominating measure μ.

We are now ready to prove Theorem 1.2.

Proof of Theorem 1.2. We break the proof into two separate cases: the case when $\inf_{\theta \in \Theta} t_\theta$ is strictly smaller than some constant b and the case when $\inf_{\theta \in \Theta} t_\theta$ is larger than b. The first case is the easy case where we show that $\widehat{\theta}(X; \Theta, f)$ is C-Bayes with respect to a simple two-point prior via Le Cam's classical two-point testing inequality. The second case is harder where we use a more elaborate prior w together with inequality (3.1) to lower bound $R_{\text{Bayes}}(w)$.

Easy Case: Here we assume that $\inf_{\theta \in \Theta} t_\theta \leq b$ (the precise value of the constant b will be specified later). Choose $\theta^* \in \Theta$ such that $t_{\theta^*} \leq b$ (note that $\theta \mapsto t_\theta$ is continuous from (2.7) and that Θ is closed so that such a θ^* exists). Let $\theta_1 \in \Theta$ be any maximizer of $\|\theta^* - \theta\|_2$ as θ varies over $\{\theta \in \Theta : \|\theta - \theta^*\|_2 \leq 1\}$. Let w be the uniform prior over the two-point set $\{\theta^*, \theta_1\}$. The Bayes risk with

respect to w can be easily bounded by below by Le Cam's inequality (from Le Cam (1973)) which gives

$$R_{\text{Bayes}}(w) \geq \frac{1}{4}\|\theta^* - \theta_1\|_2^2 \left(1 - \|P_{\theta^*} - P_{\theta_1}\|_{TV}\right)$$

where $\|P_{\theta^*} - P_{\theta_1}\|_{TV}$ denotes the total variation distance between the probability measures P_{θ^*} and P_{θ_1}. Pinsker's inequality (see for example (Tsybakov, 2009, Lemma 2.5)) now implies

$$2\|P_{\theta^*} - P_{\theta_1}\|_{TV}^2 \leq D(P_{\theta^*}\|P_{\theta_1}) = \frac{1}{2}\|\theta^* - \theta_1\|_2^2 \leq \frac{1}{2}$$

and hence

$$R_{\text{Bayes}}(w) \geq \frac{1}{8}\|\theta^* - \theta_1\|_2^2. \tag{3.3}$$

By the definition of θ_1, we have $\|\theta_1 - \theta^*\|_2 \leq 1$. We consider the following two cases by the value of $\|\theta_1 - \theta^*\|_2$.

1. $\|\theta^* - \theta_1\|_2 = 1$: Here inequality (3.3) gives $R_{\text{Bayes}}(w) \geq 1/8$. Further, by the assumption $t_{\theta^*} \leq b$ and inequality (2.3), we have

$$\mathbb{E}_{\theta^*}\|\widehat{\theta}(X) - \theta^*\|_2^2 \leq b^2 + (2\sqrt{84})b^{3/2} + 84b$$

Moreover, by inequality (2.6) and (2.3), we have

$$\begin{aligned}\mathbb{E}_{\theta_1}\|\widehat{\theta}(X;\Theta,f) - \theta_1\|_2^2 &\leq 2\mathbb{E}_{\theta^*}\|\widehat{\theta}(X;\Theta,f) - \theta^*\|_2^2 + 8\|\theta^* - \theta_1\|_2^2 \\ &\leq 2\left(b^2 + (2\sqrt{84})b^{3/2} + 84b\right) + 8\|\theta^* - \theta_1\|_2^2 \\ &\leq 2\left(b^2 + (2\sqrt{84})b^{3/2} + 84b\right) + 8.\end{aligned}$$

Combining the above two inequalities, we deduce that

$$\int_\Theta \mathbb{E}_\theta\|\widehat{\theta}(X;\Theta,f) - \theta\|_2^2 dw(\theta) \leq \frac{3}{2}\left(b^2 + (2\sqrt{84})b^{3/2} + 84b\right) + 4.$$

This inequality together with $R_{\text{Bayes}}(w) \geq 1/8$ allow us to obtain

$$\frac{1}{12\left(b^2 + (2\sqrt{84})b^{3/2} + 84b\right) + 32}\int_\Theta \mathbb{E}_\theta\|\widehat{\theta}(X;\Theta,f) - \theta\|_2^2 dw(\theta) \leq R_{\text{Bayes}}(w).$$

This means that $\widehat{\theta}(X;\Theta,f)$ is C-Bayes with respect to w with

$$C := \frac{1}{12\left(b^2 + (2\sqrt{84})b^{3/2} + 84b\right) + 32}. \tag{3.4}$$

2. $\|\theta^* - \theta_1\|_2 < 1$: In this case, $\gamma := \text{diam}(\Theta) \leq 2$ and $\|\theta^* - \theta_1\|_2 \geq \gamma/2$. Inequality (3.3) then gives $R_{\text{Bayes}}(w) \geq \gamma^2/32$. Also for every $\theta \in \Theta$, we have $\mathbb{E}_\theta\|\widehat{\theta}(X;\Theta,f) - \theta\|_2^2 \leq \gamma^2$ (because both $\widehat{\theta}(X;\Theta,f)$ and θ are

constrained to take values in Θ whose diameter is at most γ). These two inequalities imply that

$$\frac{1}{32}\int_\Theta \mathbb{E}_\theta \|\widehat{\theta}(X;\Theta,f)-\theta\|_2^2 dw(\theta) \le R_{\text{Bayes}}(w)$$

which means that $\widehat{\theta}(X;\Theta,f)$ is C-Bayes with respect to w with $C = 1/32$.

Therefore in this easy case, we have proved that $\widehat{\theta}(X;\Theta,f)$ is C-Bayes for some C that is atleast the minimum of (3.4) and $1/32$.

Hard Case: We now work with the situation when $\inf_{\theta\in\Theta} t_\theta > b$. We fix a specific $\theta^* \in \Theta$ and choose w as a specific prior that is supported on the set

$$U(\theta^*) := \Theta \cap \{\theta \in \mathbb{R}^n : \|\theta-\theta^*\|_2 \le \rho t_{\theta^*}\} \tag{3.5}$$

for some constant $\rho > 0$ (to be specified later) which satisfies $\rho^2 + 4\rho < 1$. More precisely, for a fixed small constant η, let θ^* be chosen so that

$$m_{\theta^*}(\rho t_{\theta^*}) \ge \sup_{\theta\in\Theta} m_\theta(\rho t_\theta) - \eta \tag{3.6}$$

where $m_\theta(\cdot)$ is defined in (2.1). Let $\Psi : \mathbb{R}^n \mapsto \Theta$ be any measurable mapping such that $\Psi(z)$ is a maximizer of $\langle z, \alpha-\theta^*\rangle - f(\alpha)$ as α varies in $U(\theta^*)$. Let w be the prior given by the distribution of $\Psi(Z)$ for a standard Gaussian vector Z in $\mathbb{R}^n$.

Now because of inequalities (2.3) and (2.6), we can write the following for every $\theta \in U(\theta^*)$:

$$\begin{aligned}
\mathbb{E}_\theta\|\widehat{\theta}(X;\Theta,f)-\theta\|_2^2 &\le 2\mathbb{E}_{\theta^*}\|\widehat{\theta}(X;\Theta,f)-\theta^*\|_2^2 + 8\|\theta-\theta^*\|_2^2 \\
&\le 2\left(t_{\theta^*}^2 + (2\sqrt{84})t_{\theta^*}^{3/2} + 84 t_{\theta^*}\right) + 8\|\theta^*-\theta\|_2^2 \\
&\le 2\left(t_{\theta^*}^2 + (2\sqrt{84})t_{\theta^*}^{3/2} + 84 t_{\theta^*}\right) + 8\rho^2 t_{\theta^*}^2 \\
&= (2+8\rho^2)t_{\theta^*}^2 + (4\sqrt{84})t_{\theta^*}^{3/2} + 168 t_{\theta^*}.
\end{aligned}$$

This clearly implies

$$\begin{aligned}
\int_\Theta \mathbb{E}_\theta\|\widehat{\theta}(X;\Theta,f)-\theta\|_2^2 dw(\theta) &\le (2+8\rho^2)t_{\theta^*}^2 + (4\sqrt{84})t_{\theta^*}^{3/2} + 168 t_{\theta^*} \\
&\le t_{\theta^*}^2\left(2+8\rho^2+4\sqrt{84}b^{-1/2}+168 b^{-1}\right) \qquad (3.7)
\end{aligned}$$

where the second inequality above follows from the fact that $t_{\theta^*} \ge \inf_{\theta\in\Theta} t_\theta > b$.

The goal now is to provide a lower bound for $R_{\text{Bayes}}(w)$. We shall use inequality (3.1) for this purpose. Because the prior w is concentrated on the convex set $U(\theta^*)$, we can replace the supremum over $a \in \Theta$ in (3.1) by the supremum over $a \in U(\theta^*)$. This gives the following lower bound for $R_{\text{Bayes}}(w)$:

$$R_{\text{Bayes}}(w) \ge \frac{1}{2}\sup\left\{t>0 : \sup_{a\in U(\theta^*)} w\{\theta\in\Theta : \|\theta-a\|_2^2 \le t\} < \frac{1}{4(1+I)}\right\} \tag{3.8}$$

where I is any upper bound on $\inf_Q \int_\Theta \chi^2(P_\theta \| Q) dw(\theta)$. Here P_θ is the n-dimensional normal distribution with mean zero and identity covariance and the infimum is over all probability measures Q.

To obtain a suitable value for I, we use

$$\begin{aligned}\inf_Q \int_\Theta \chi^2(P_\theta \| Q) dw(\theta) &\leq \int_\Theta \chi^2(P_\theta \| P_{\theta^*}) dw(\theta) \\ &= \int_{U(\theta^*)} \chi^2(P_\theta \| P_{\theta^*}) dw(\theta) \\ &\leq \sup_{\theta \in U(\theta^*)} \chi^2(P_\theta \| P_{\theta^*}) \leq \exp(\rho^2 t_{\theta^*}^2) - 1\end{aligned}$$

where, in the last inequality, we used the expression $\chi^2(P_\theta \| P_{\theta^*}) = \exp(\|\theta - \theta^*\|_2^2) - 1$ and the fact that $\|\theta - \theta^*\|_2 \leq \rho t_{\theta^*}$ for all $\theta \in U(\theta^*)$. We can therefore take $1 + I$ to be $\exp(\rho^2 t_{\theta^*}^2)$ in (3.8) which gives

$$R_{\text{Bayes}}(w) \geq \frac{1}{2} \sup \left\{ t > 0 : \sup_{a \in U(\theta^*)} w\{\theta \in \Theta : \|\theta - a\|_2^2 \leq t\} < \frac{1}{4} \exp(-\rho^2 t_{\theta^*}^2) \right\}. \tag{3.9}$$

We shall now bound from above

$$\sup_{a \in U(\theta^*)} w\{\theta \in \Theta : \|\theta - a\|_2^2 \leq t\} \qquad \text{for } \sqrt{t} = (1-\beta)\rho \left(1 - \sqrt{\rho^2 + 4\rho}\right) t_{\theta^*} \tag{3.10}$$

for a constant $\beta \in (0,1)$. The goal is to show that the above quantity is smaller than $\exp(-\rho^2 t_{\theta^*}^2)/4$.

Because w is defined as the distribution of $\Psi(Z)$ which is a maximizer of $\langle Z, \alpha - \theta^* \rangle - f(\alpha)$ over $\alpha \in U(\theta^*)$, the inequality

$$w(A) \leq \mathbb{P} \left\{ \sup_{\alpha \in A} (\langle Z, \alpha - \theta^* \rangle - f(\alpha)) \geq \sup_{\alpha \in U(\theta^*)} (\langle Z, \alpha - \theta^* \rangle - f(\alpha)) \right\}$$

holds for every measureable subset A of $\mathbb{R}^n$. Therefore for every $a \in U(\theta^*)$, the prior probability $w\{\theta \in \Theta : \|\theta - a\|_2^2 \leq t\}$ is bounded from above by

$$\mathbb{P} \left\{ \sup_{\theta \in \Theta : \|\theta - a\|_2^2 \leq t} (\langle Z, \theta - \theta^* \rangle - f(\theta)) \geq \sup_{\theta \in \Theta : \|\theta - \theta^*\| \leq \rho t_{\theta^*}} (\langle Z, \theta - \theta^* \rangle - f(\theta)) \right\}.$$

The above probability can be exactly written as $\mathbb{P}\{M_2 + M_3 \geq M_1\}$ where

$$M_1 := \sup_{\theta \in \Theta : \|\theta - \theta^*\| \leq \rho t_{\theta^*}} (\langle Z, \theta - \theta^* \rangle - f(\theta)), \quad M_2 := \sup_{\theta \in \Theta : \|\theta - a\|_2^2 \leq t} (\langle Z, \theta - a \rangle - f(\theta))$$

and $M_3 := \langle Z, a - \theta^* \rangle$. Now if $\gamma \geq 0$ is such that $\mathbb{E}M_1 - \mathbb{E}M_2 \geq \gamma$, then we can

write:

$$\begin{aligned}\mathbb{P}\{M_2 + M_3 \geq M_1\} &= \mathbb{P}\left\{M_2 - \mathbb{E}M_2 + M_3 + \mathbb{E}M_1 - M_1 \geq \mathbb{E}M_1 - \mathbb{E}M_2\right\} \\ &\leq \mathbb{P}\left\{M_2 - \mathbb{E}M_2 + M_3 + \mathbb{E}M_1 - M_1 \geq \gamma\right\} \\ &\leq \mathbb{P}\left\{M_2 - \mathbb{E}M_2 \geq \frac{\gamma}{3}\right\} + \mathbb{P}\left\{M_3 \geq \frac{\gamma}{3}\right\} + \mathbb{P}\left\{M_1 - \mathbb{E}M_1 \leq \frac{-\gamma}{3}\right\} \\ &\leq \exp\left(\frac{-\gamma^2}{18t}\right) + \exp\left(\frac{-\gamma^2}{18\|a-\theta^*\|_2^2}\right) + \exp\left(\frac{-\gamma^2}{18\rho^2 t_{\theta^*}^2}\right).\end{aligned}$$

where the last inequality follows by standard Gaussian concentration and the observation that (a) M_2, as a function of Z, is Lipschitz with Lipschitz constant $\sqrt{t}$, (b) M_3, as a function of Z, is Lipschitz with Lipschitz constant $\|a-\theta^*\|_2$, and (c) M_1, as a function of Z, is Lipschitz with Lipschitz constant ρt_{θ^*}.

We now use the fact that for every $a \in U(\theta^*)$, the inequality $\|a-\theta^*\|_2 \leq \rho t_{\theta^*}$ holds to deduce that

$$\begin{aligned}w\left\{\theta \in \Theta : \|\theta - a\|_2^2 \leq t\right\} &\leq \exp\left(\frac{-\gamma^2}{18t}\right) + 2\exp\left(\frac{-\gamma^2}{18\rho^2 t_{\theta^*}^2}\right) \\ &\leq 3\exp\left(\frac{-\gamma^2}{18\rho^2 t_{\theta^*}^2}\right) \quad \text{as } t \leq \rho^2 t_{\theta^*}^2 (\text{see (3.10)}) \quad (3.11)\end{aligned}$$

Here γ is any nonnegative lower bound on $\mathbb{E}M_1 - \mathbb{E}M_2$. To obtain a suitable value of γ, we argue as follows. Observe first that $\mathbb{E}M_1 = m_{\theta^*}(\rho t_{\theta^*})$ and $\mathbb{E}M_2 = m_a(\sqrt{t})$ where m is defined in (2.1). Because θ^* is chosen so that inequality (3.6) is satisfied, we have $\mathbb{E}M_1 = m_{\theta^*}(\rho t_{\theta^*}) \geq m_a(\rho t_a) - \eta$. Thus

$$\mathbb{E}M_1 - \mathbb{E}M_2 \geq m_a(\rho t_a) - m_a(\sqrt{t}) - \eta.$$

We now use inequality (2.8) which states that

$$t_a \geq \left(1 - \sqrt{\rho^2 + 4\rho}\right) t_{\theta^*} \qquad \text{for every } a \in U(\theta^*). \qquad (3.12)$$

Now from the expression for t given in (3.10) and inequality (3.12) above, it is clear that $\sqrt{t} \leq \rho t_a$ for every $a \in U(\theta^*)$. Therefore using concavity of $m_a(\cdot)$ (proved in Lemma 2.2) and inequality (2.4), we deduce

$$\begin{aligned}\mathbb{E}M_1 - \mathbb{E}M_2 &\geq m_a(\rho t_a) - m_a(\sqrt{t}) - \eta \\ &\geq m_a(t_a) - m_a(t_a - \rho t_a + \sqrt{t}) - \eta \qquad \text{by concavity of } m_a(\cdot) \\ &\geq t_a\left(\rho t_a - \sqrt{t}\right) - \eta \qquad \text{by inequality (2.4)} \\ &\geq \rho\beta t_{\theta^*}^2 \left(1 - \sqrt{\rho^2 + 4\rho}\right)^2 - \eta \qquad \text{by inequality (3.12) and the expression for } t.\end{aligned}$$

We therefore take γ to be

$$\gamma = \rho\beta t_{\theta^*}^2 \left(1 - \sqrt{\rho^2 + 4\rho}\right)^2 - \eta.$$

Inequality (3.11) then gives

$$\sup_{a\in U(\theta^*)} w\left\{\theta\in\Theta : \|\theta-a\|_2^2 \le t\right\} \le 3\exp\left(\frac{-\left\{\rho\beta t_{\theta^*}(1-\sqrt{\rho^2+4\rho})^2-\eta\right\}^2}{18\rho^2 t_{\theta^*}^2}\right).$$

By a straightforward computation, it can be seen that the right hand side above is strictly smaller than $\frac{1}{4}\exp(-\rho^2 t_{\theta^*}^2)$ if and only if

$$\frac{1}{18\rho^2 t_{\theta^*}^2}\left(\rho\beta t_{\theta^*}^2\left(1-\sqrt{\rho^2+4\rho}\right)^2-\eta\right)^2-\rho^2 t_{\theta^*}^2 > \log 12 \tag{3.13}$$

Now, as a result of the following inequality (note that we are working under the condition $\inf_{\theta\in\Theta} t_\theta > b$ which implies that $t_{\theta^*} > b$):

$$\begin{aligned}\frac{1}{18\rho^2 t_{\theta^*}^2}\left(\rho\beta t_{\theta^*}^2\left(1-\sqrt{\rho^2+4\rho}\right)^2-\eta\right)^2 &\ge \frac{1}{18\rho^2 t_{\theta^*}^2}\left(\rho\beta t_{\theta^*}^2\left(1-\sqrt{\rho^2+4\rho}\right)^2-\eta\frac{t_{\theta^*}^2}{b^2}\right)^2\\ &= \frac{t_{\theta^*}^2}{18\rho^2}\left(\rho\beta\left(1-\sqrt{\rho^2+4\rho}\right)^2-\eta b^{-2}\right)^2,\end{aligned}$$

a sufficient condition for (3.13) is

$$t_{\theta^*}^2 > (\log 12)\left(\frac{1}{18\rho^2}\left(\rho\beta(1-\sqrt{\rho^2+4\rho})^2-\eta b^{-2}\right)^2-\rho^2\right)^{-1} \tag{3.14}$$

We now make the choices:

$$\rho = 0.0295 \qquad \beta = 0.42 \qquad \eta = 10^{-20} \qquad b = 51.53. \tag{3.15}$$

With these, the right hand side of (3.14) can be calculated to be strictly smaller than b^2 so that the condition (3.14) holds because $t_{\theta^*} > b$. Therefore we deduce from inequality (3.9) that

$$R_{\text{Bayes}}(w) \ge \frac{t}{2} = \frac{\rho^2}{2}(1-\beta)^2\left(1-\sqrt{\rho^2+4\rho}\right)^2 t_{\theta^*}^2.$$

Combining the above inequality with (3.7), we obtain

$$\frac{\rho^2(1-\beta)^2(1-\sqrt{\rho^2+4\rho})^2}{2(2+8\rho^2+4\sqrt{84}b^{-1/2}+168b^{-1})}\int_\Theta \mathbb{E}_\theta\|\widehat{\theta}(X;\Theta,f)-\theta\|_2^2 dw(\theta) \le R_{\text{Bayes}}(w)$$

The constant above (for our choice of ρ, β and b in (3.15)) is at least 6.05×10^{-6}. This means therefore that $\widehat{\theta}(X;\Theta,f)$ is C-Bayes with respect to w with $C \ge 6.05\times 10^{-6}$ in the case when $\inf_{\theta\in\Theta} t_\theta > b = 51.53$.

It is also easy to check that for $b = 51.53$, the constant in (3.4) is also at least 6.05×10^{-6}. Therefore in every case, we have proved the existence of a prior w such that $\widehat{\theta}(X;\Theta,f)$ is C-Bayes with respect to w for $C \ge 6.05\times 10^{-6}$. This means that $\widehat{\theta}(X;\Theta,f)$ is C-admissible for some constant $C \ge 6.05\times 10^{-6}$. This completes the proof of Theorem 1.2. □

4. Proof of Theorem 1.3

Theorem 1.2 shows that for every $n \geq 1$, $\Theta \in \mathcal{C}_n$ and $f \in \mathcal{F}(\Theta)$, the estimator $\widehat{\theta}(X; \Theta, f)$ is C-admissible over Θ for some $C \geq 6.05 \times 10^{-6}$. This immediately implies that $C^* \geq 6.05 \times 10^{-6}$. We therefore only need to prove that $C^* \leq 1/2$.

For this it is enough to show that for every $\epsilon > 0$, there exists a closed convex set $\Theta \subseteq \mathbb{R}$ such that the LSE $\widehat{\theta}(X; \Theta)$ satisfies

$$\inf_{\mathfrak{d}} \sup_{\theta \in \Theta} \frac{\mathbb{E}_\theta \|\mathfrak{d}(X) - \theta\|_2^2}{\mathbb{E}_\theta \|\widehat{\theta}(X; \Theta) - \theta\|_2^2} \leq \frac{1}{2} + \epsilon. \tag{4.1}$$

Note that $\widehat{\theta}(X; \Theta)$ is a special case of $\widehat{\theta}(X; \Theta, f)$ for $f \equiv 0$.

Let $\Theta := [-a, a]$ for some $a > 0$ (to be specified later). It is then clear that

$$\widehat{\theta}(X; \Theta) = \begin{cases} X & \text{if } -a \leq X \leq a \\ a & \text{if } X > a \\ -a & \text{if } X < -a \end{cases}$$

Note now that

$$\inf_{\theta \in \Theta} \mathbb{P}_\theta\{X > a\} = P_{-a}\{X > a\} = 1 - \Phi(2a).$$

and similarly

$$\inf_{\theta \in \Theta} \mathbb{P}_\theta\{X < -a\} = 1 - \Phi(2a).$$

Therefore for every $\theta \in \Theta = [-a, a]$, we have

$$\begin{aligned} \mathbb{E}_\theta \|\widehat{\theta}(X; \Theta) - \theta\|_2^2 &\geq \mathbb{E}_\theta \left(\|\widehat{\theta}(X; \Theta) - \theta\|_2^2 \{X > a\} \right) + \mathbb{E}_\theta \left(\|\widehat{\theta}(X; \Theta) - \theta\|_2^2 \{X < -a\} \right) \\ &= (\theta - a)^2 \mathbb{P}_\theta\{X > a\} + (\theta + a)^2 \mathbb{P}_\theta\{X < -a\} \\ &\geq (\theta - a)^2 \mathbb{P}_{-a}\{X > a\} + (\theta + a)^2 \mathbb{P}_a\{X < -a\} \\ &\geq (1 - \Phi(2a)) \left[(\theta - a)^2 + (\theta + a)^2 \right] = 2 \left(1 - \Phi(2a)\right) \left(\theta^2 + a^2\right) \end{aligned}$$

where in the third line above, we used the fact that $\inf_{\theta \in [-a,a]} \mathbb{P}_\theta\{X > a\} = \mathbb{P}_{-a}\{X > a\}$ and $\inf_{\theta \in [-a,a]} \mathbb{P}_\theta\{X < -a\} = \mathbb{P}_a\{X > a\}$

As a result, we deduce that

$$\begin{aligned} \inf_{\mathfrak{d}} \sup_{\theta \in \Theta} \frac{\mathbb{E}_\theta \|\mathfrak{d}(X) - \theta\|_2^2}{\mathbb{E}_\theta \|\widehat{\theta}(X; \Theta) - \theta\|_2^2} &\leq \sup_{\theta \in \Theta} \frac{\theta^2}{\mathbb{E}_\theta \|\widehat{\theta}(X; \Theta) - \theta\|_2^2} \\ &\leq \frac{1}{2(1 - \Phi(2a))} \sup_{\theta \in \Theta} \frac{\theta^2}{\theta^2 + a^2} = \frac{1}{4(1 - \Phi(2a))} \end{aligned}$$

where in the first inequality above, we bounded the infimum over all estimators $\mathfrak{d}$ by the simple estimator $\mathfrak{d}(X) \equiv 0$. Note now that by taking the limit $a \downarrow 0$, the right hand side above goes to $1/2$. Therefore it is possible to choose a small enough depending on ϵ to ensure (4.1). This proves $C^* \leq 1/2$ and completes the proof of Theorem 1.3.

5. Proofs of Theorem 2.1 and Lemma 2.2

In this section, we shall provide proofs for Theorem 2.1 and Lemma 2.2. We first give the proof of Lemma 2.2 below before proceeding to the proof of Theorem 2.1. This is because parts of Lemma 2.2 will be useful for proving Theorem 2.1.

Proof of Lemma 2.2 . **1.** Fix $\theta \in \Theta$. We need to prove that $m_\theta(\cdot)$ is a non-decreasing and concave function. It is trivial to see that m_θ is a non-decreasing function because the sets $\{\alpha \in \Theta : \|\alpha - \theta\|_2 \le t\}$ are increasing in t. To prove concavity of m_θ, observe that it is enough to show that

$$H_\theta(t) := \sup_{\alpha \in \Theta : \|\alpha - \theta\|_2 \le t} (\langle X - \theta, \alpha - \theta \rangle - f(\alpha))$$

is concave for every $z \in \mathbb{R}^n$. This is because $m_\theta(t) = \mathbb{E} H_\theta(t)$. To prove concavity of H_θ, let us fix $0 \le t_1 < t_2 < \infty$, $0 < u < 1$ and $t = (1-u)t_1 + ut_2$. For every $\eta > 0$, by the definition of $H_\theta(t)$, for each $i = 1, 2$, there exists $\theta_i \in \Theta$ with $\|\theta_i - \theta\| \le t_i$ such that

$$H_\theta(t_i) \le \langle X - \theta, \theta_i - \theta \rangle - f(\theta_i) + \eta.$$

Now with $\alpha := (1-u)\theta_1 + u\theta_2$, it is easy to see that $\|\alpha - \theta\|_2 \le t, \alpha \in \Theta$ and, by convexity of f, that

$$\begin{aligned} H_\theta(t) &\ge \langle X - \theta, \alpha - \theta \rangle - f(\alpha) \\ &\ge (1-u)\left[\langle X - \theta, \theta_1 - \theta \rangle - f(\theta_1)\right] + u\left[\langle X - \theta, \theta_2 - \theta \rangle - f(\theta_2)\right] \\ &\ge (1-u) H_\theta(t_1) + u H_\theta(t_2) - \eta. \end{aligned}$$

Because $\eta > 0$ is arbitrary, this proves concavity of H_θ on $[0, \infty)$ which implies concavity of m_θ on $[0, \infty)$.

2. The concavity of m_θ implies that $G_\theta(t) := m_\theta(t) - t^2/2$ is strictly concave on $[0, \infty)$. Moreover

$$\begin{aligned} m_\theta(t) &= \mathbb{E}_\theta \left(\sup_{\alpha \in \Theta : \|\alpha - \theta\|_2 \le t} \{\langle X - \theta, \alpha - \theta \rangle - f(\alpha)\} \right) \\ &\le \mathbb{E}_\theta \left(\sup_{\alpha \in \Theta : \|\alpha - \theta\|_2 \le t} \langle X - \theta, \alpha - \theta \rangle \right) - \inf_{\alpha \in \Theta} f(\alpha) \\ &\le \|X - \theta\|_2 t - \inf_{\alpha \in \Theta} f(\alpha) \end{aligned}$$

where, in the last inequality above, we used the Cauchy-Schwarz inequality. As a result, $G_\theta(t) = m_\theta(t) - t^2/2$ converges to $-\infty$ as $t \uparrow +\infty$. This, together with strict concavity of $G_\theta(\cdot)$ on $[0, \infty)$, immediately imply that G_θ has a unique maximizer t_θ over $[0, \infty)$.

3. Fix $\theta \in \Theta$ and $t \ge 0$. Then the inequality $G_\theta(t_\theta) \ge G_\theta(t)$ holds because t_θ maximizes $G_\theta(\cdot)$. This inequality is equivalent to

$$m_\theta(t) \le m_\theta(t_\theta) + \frac{t^2 - t_\theta^2}{2}.$$

Applying this inequality to $(1-u)t_\theta + ut$ instead of t for a fixed $u \in (0,1)$, we obtain

$$m_\theta((1-u)t_\theta + ut) \le m_\theta(t_\theta) + \frac{(-2u+u^2)t_\theta^2 + u^2t^2 + 2u(1-u)t\ t_\theta}{2}.$$

Using $m_\theta((1-u)t_\theta + ut) \ge (1-u)m_\theta(t_\theta) + u\ m_\theta(t)$ on the right hand side above, we get

$$u\ m_\theta(t) \le u\ m_\theta(t_\theta) + \frac{(-2u+u^2)t_\theta^2 + u^2t^2 + 2u(1-u)t\ t_\theta}{2}..$$

Dividing both sides of the above inequality by u and then letting $u \to 0$, we obtain

$$m_\theta(t_\theta) - m_\theta(t) \ge t_\theta^2 - t\ t_\theta$$

which proves (2.4).

4. The expression $G_\theta(t) := m_\theta(t) - t^2/2$ implies that inequality (2.5) is equivalent to (2.4). Therefore, inequality (2.5) follows from the previous part.

5. To prove (2.6), we first observe that the map $X \mapsto \widehat{\theta}(X;\Theta,f)$ is 1-Lipschitz (for a proof of this standard fact, see e.g., van de Geer and Wainwright (2015, Proof of Theorem 2.1)). As a result, for every $\theta_1, \theta_2 \in \Theta$, we can write (for simplicity below, we write $\widehat{\theta}(X)$ for $\widehat{\theta}(X;\Theta,f)$)

$$\begin{aligned}\|\widehat{\theta}(X) - \theta_1\|_2 &\le \|\widehat{\theta}(X) - \widehat{\theta}(X-\theta_1+\theta_2)\|_2 + \|\widehat{\theta}(X-\theta_1+\theta_2) - \theta_2\|_2 + \|\theta_1-\theta_2\|_2 \\ &\le \|\widehat{\theta}(X-\theta_1+\theta_2) - \theta_2\|_2 + 2\|\theta_1-\theta_2\|_2\end{aligned}$$

where the first inequality is due to the triangle inequality while the second inequality is because $\|\widehat{\theta}(X) - \widehat{\theta}(X-\theta_1+\theta_2)\|_2 \le \|\theta_1-\theta_2\|_2$ by the 1-Lipschitz property of $X \mapsto \widehat{\theta}(X)$. By squaring both sides of the above displayed inequality and using $(a+b)^2 \le 2a^2 + 2b^2$, we get

$$\|\widehat{\theta}(X) - \theta_1\|_2^2 \le 2\|\widehat{\theta}(X-\theta_1+\theta_2) - \theta_2\|_2^2 + 8\|\theta_1-\theta_2\|_2^2.$$

Taking expectations with respect to $X \sim N(\theta_1, I_n)$ on both sides and using the fact that $X - \theta_1 + \theta_2 \sim N(\theta_2, I_n)$, we obtain the required inequality (2.6).

6. Fix $\theta_1, \theta_2 \in \Theta$ and observe that, for every $t \ge 0$, we have

$$\begin{aligned}m_{\theta_2}(t) &= \mathbb{E}_{\theta_2}\left(\sup_{\alpha\in\Theta:\|\alpha-\theta_2\|_2\le t}\{\langle X-\theta_2, \alpha-\theta_2\rangle - f(\alpha)\}\right) \\ &= \mathbb{E}_{\theta_2}\left(\sup_{\alpha\in\Theta:\|\alpha-\theta_2\|_2\le t}\{\langle X-\theta_2, \alpha\rangle - f(\alpha)\}\right) \\ &\le \mathbb{E}_{\theta_2}\left(\sup_{\alpha\in\Theta:\|\alpha-\theta_1\|_2\le t+\|\theta_1-\theta_2\|_2}\{\langle X-\theta_2, \alpha\rangle - f(\alpha)\}\right) \\ &= \mathbb{E}_{\theta_1}\left(\sup_{\alpha\in\Theta:\|\alpha-\theta_1\|_2\le t+\|\theta_1-\theta_2\|_2}\{\langle X-\theta_1, \alpha-\theta_1\rangle - f(\alpha)\}\right) = m_{\theta_1}(t+\|\theta_1-\theta_2\|_2).\end{aligned}$$

Switching the roles of θ_1 and θ_2, we obtain

$$m_{\theta_1}(t) \leq m_{\theta_2}(t + \|\theta_1 - \theta_2\|_2).$$

Combining the above two inequalities, we deduce that

$$m_{\theta_1}(t - \|\theta_1 - \theta_2\|_2) \leq m_{\theta_2}(t) \leq m_{\theta_1}(t + \|\theta_1 - \theta_2\|_2)$$

provided $\|\theta_1 - \theta_2\|_2 \leq t$. Using $G_{\theta_2}(t) = m_{\theta_2}(t) - t^2/2$, we further deduce that

$$m_{\theta_1}(t - \|\theta_1 - \theta_2\|_2) - \frac{t^2}{2} \leq G_{\theta_2}(t) \leq m_{\theta_1}(t + \|\theta_1 - \theta_2\|_2) - \frac{t^2}{2}$$

Therefore for every $t \geq 0$, we have

$$\begin{aligned} G_{\theta_2}(t_{\theta_1} + \|\theta_1 - \theta_2\|_2) - G_{\theta_2}(t) &\geq \left\{ m_{\theta_1}(t_{\theta_1}) - \frac{1}{2}\left(t_{\theta_1} + \|\theta_1 - \theta_2\|_2\right)^2 \right\} - \left\{ m_{\theta_1}(t + \|\theta_1 - \theta_2\|_2) - \frac{t^2}{2} \right\} \\ &= \left\{ m_{\theta_1}(t_{\theta_1}) - m_{\theta_1}(t + \|\theta_1 - \theta_2\|_2) \right\} - \frac{1}{2}\left\{ \left(t_{\theta_1} + \|\theta_1 - \theta_2\|_2\right)^2 - t^2 \right\} \\ &\geq t_{\theta_1}\left(t_{\theta_1} - t - \|\theta_1 - \theta_2\|_2\right) - \frac{1}{2}\left\{ \left(t_{\theta_1} + \|\theta_1 - \theta_2\|_2\right)^2 - t^2 \right\} \end{aligned}$$

where we used (2.4) for the last inequality above. Simplifying the right hand side above, we deduce that

$$G_{\theta_2}(t_{\theta_1} + \|\theta_1 - \theta_2\|_2) - G_{\theta_2}(t) \geq \frac{1}{2}\left(t - t_{\theta_1}\right)^2 - \frac{1}{2}\|\theta_1 - \theta_2\|_2^2 - 2t_{\theta_1}\|\theta_1 - \theta_2\|_2. \quad (5.1)$$

From this expression, it follows that

$$G_{\theta_2}(t_{\text{low}}) \leq G_{\theta_2}(t_{\theta_1} + \|\theta_1 - \theta_2\|_2) \leq G_{\theta_2}(t_{\text{up}}) \quad (5.2)$$

for

$$t_{\text{low}} := t_{\theta_1} - \sqrt{\|\theta_1 - \theta_2\|_2^2 + 4t_{\theta_1}\|\theta_1 - \theta_2\|_2} \quad \text{and} \quad t_{\text{up}} := t_{\theta_1} + \sqrt{\|\theta_1 - \theta_2\|_2^2 + 4t_{\theta_1}\|\theta_1 - \theta_2\|_2}$$

as long as $t_{\text{low}} \geq 0$. Inequality (5.2), together with the fact that $G_{\theta_2}(\cdot)$ is strictly concave on $[0, \infty)$ (this follows from concavity of $m_\theta(\cdot)$ and the observation that $t \mapsto -t^2/2$ is strictly concave) implies that $t_{\text{low}} \leq t_{\theta_2} \leq t_{\text{up}}$ as long as $t_{\text{low}} \geq 0$. Because t_{θ_2} is always nonnegative, we deduce therefore that

$$\left(t_{\theta_1} - \sqrt{\|\theta_1 - \theta_2\|_2^2 + 4t_{\theta_1}\|\theta_1 - \theta_2\|_2}\right)_+ \leq t_{\theta_2} \leq t_{\theta_1} + \sqrt{\|\theta_1 - \theta_2\|_2^2 + 4t_{\theta_1}\|\theta_1 - \theta_2\|_2}$$

for every $\theta_1, \theta_2 \in \Theta$ which proves (2.7). It is easy to see that inequality (2.8) is a simple consequence of (2.7). □

Proof of Theorem 2.1. Let $Z := X - \theta \sim N(0, I_n)$. By the calculation

$$\|X - \alpha\|_2^2 = \|Z + \theta - \alpha\|_2^2 = \|Z\|_2^2 - 2\langle Z, \alpha - \theta\rangle + \|\alpha - \theta\|_2^2,$$

it follows that

$$\widehat{\theta}(X;\Theta,f) = \operatorname*{argmax}_{\alpha\in\Theta}\left\{\langle Z,\alpha-\theta\rangle - f(\alpha) - \frac{1}{2}\|\alpha-\theta\|_2^2\right\} \tag{5.3}$$

For convenience, we shall write $\widehat{\theta}$ for $\widehat{\theta}(X;\Theta,f)$ in the rest of the proof.

From (5.3), one can see as follows that $\|\widehat{\theta}-\theta\|_2$ maximizes the function $\tilde{A}_\theta(t)$ over $t \geq 0$ where

$$\tilde{A}_\theta(t) := \sup_{\theta\in\Theta:\|\alpha-\theta\|_2=t}\{\langle Z,\alpha-\theta\rangle - f(\alpha)\} - \frac{t^2}{2}.$$

To see this, just observe that

$$\begin{aligned}
\tilde{A}_\theta(\|\widehat{\theta}-\theta\|_2) &= \sup_{\theta\in\Theta:\|\alpha-\theta\|_2=\|\widehat{\theta}-\theta\|_2}\{\langle Z,\alpha-\theta\rangle - f(\alpha)\} - \frac{\|\widehat{\theta}-\theta\|_2^2}{2} \\
&\geq \left\langle Z,\widehat{\theta}-\theta\right\rangle - f(\widehat{\theta}) - \frac{\|\widehat{\theta}-\theta\|_2^2}{2} \\
&= \sup_{\alpha\in\Theta}\left(\langle Z,\alpha-\theta\rangle - f(\alpha) - \frac{\|\alpha-\theta\|_2^2}{2}\right) \\
&\geq \sup_{\alpha\in\Theta:\|\alpha-\theta\|_2=t}\left(\langle Z,\alpha-\theta\rangle - f(\alpha) - \frac{\|\alpha-\theta\|_2^2}{2}\right) \qquad \text{for every } t\geq 0 \\
&= \tilde{A}_\theta(t).
\end{aligned}$$

Using this fact, we shall now argue that $\|\widehat{\theta}-\theta\|_2$ also maximizes the function $A_\theta(t)$ over $t \geq 0$ where

$$A_\theta(t) := \sup_{\theta\in\Theta:\|\alpha-\theta\|_2\leq t}\{\langle Z,\alpha-\theta\rangle - f(\alpha)\} - \frac{t^2}{2}.$$

Note that the difference between $\tilde{A}_\theta(t)$ and $A_\theta(t)$ is that the supremum is taken over $\|\alpha-\theta\|_2 = t$ in $\tilde{A}_\theta(t)$ while it is over the larger set $\|\alpha-\theta\|_2 \leq t$ in $A_\theta(t)$. This in particular means that $\tilde{A}_\theta(t) \leq A_\theta(t)$ for every $t \geq 0$. To see that $\|\widehat{\theta}-\theta\|_2$ maximizes $A_\theta(t)$, fix $t \geq 0$ and $\eta > 0$. By definition of $A_\theta(t)$, there exists $\alpha \in \Theta$ with $\|\alpha-\theta\|_2 \leq t$ such that

$$A_\theta(t) \leq \langle Z,\alpha-\theta\rangle - f(\alpha) + \eta - \frac{t^2}{2}.$$

Because $\|\alpha-\theta\|_2 \leq t$, we can write

$$\begin{aligned}
A_\theta(t) &\leq \langle Z,\alpha-\theta\rangle - f(\alpha) + \eta - \frac{t^2}{2} \\
&\leq \langle Z,\alpha-\theta\rangle - f(\alpha) + \eta - \frac{\|\alpha-\theta\|_2^2}{2} \\
&\leq \tilde{A}_\theta(\|\alpha-\theta\|_2) + \eta \\
&\leq \tilde{A}_\theta(\|\widehat{\theta}-\theta\|_2) + \eta \qquad \text{because } \|\widehat{\theta}-\theta\|_2 \text{ maximizes } \tilde{A}_\theta(\cdot) \\
&\leq A_\theta(\|\widehat{\theta}-\theta\|_2) + \eta \qquad \text{because } \tilde{A}_\theta(t) \leq A_\theta(t) \text{ for every } t \geq 0.
\end{aligned}$$

Because $t \geq 0$ and $\eta > 0$ are arbitrary, we have proved that $\|\widehat{\theta} - \theta\|_2$ maximizes $A_\theta(t)$ over $t \geq 0$.

Note now that $A_\theta(t)$ is a concave function of $t \geq 0$. This is because the function

$$t \mapsto \sup_{\alpha \in \Theta : \|\alpha - \theta\|_2 \leq t} (\langle Z, \alpha - \theta \rangle - f(\alpha))$$

is concave as shown in the proof of Lemma 2.2(1) and also $t \mapsto -t^2/2$ is trivially concave. As a result of the concavity of $A_\theta(\cdot)$, it follows that for every $\delta \geq 0$,

$$\|\widehat{\theta} - \theta\|_2 < t_\theta + \delta \tag{5.4}$$

provided

$$\begin{aligned} & \max\left(\mathbb{E}_\theta A_\theta(t_\theta) - A_\theta(t_\theta), A_\theta(t_\theta + \delta) - \mathbb{E}_\theta A_\theta(t_\theta + \delta)\right) \\ < \; & \frac{1}{2}\left\{\mathbb{E}_\theta A_\theta(t_\theta) - \mathbb{E}_\theta A_\theta(t_\theta + \delta)\right\}. \end{aligned} \tag{5.5}$$

To see this, assume that (5.5) holds for some $\delta \geq 0$. Then, if $B := (\mathbb{E}_\theta A_\theta(t_\theta) + \mathbb{E}_\theta A_\theta(t_\theta + \delta))/2$, then (5.5) implies that

$$A_\theta(t_\theta) > \mathbb{E}_\theta A_\theta(t_\theta) - \frac{1}{2}\left(\mathbb{E}_\theta A_\theta(t_\theta) - \mathbb{E}_\theta A_\theta(t_\theta + \delta)\right) = B$$

and also

$$A_\theta(t_\theta + \delta) < \mathbb{E}_\theta A_\theta(t_\theta + \delta) + \frac{1}{2}\left(\mathbb{E}_\theta A_\theta(t_\theta) - \mathbb{E}_\theta A_\theta(t_\theta + \delta)\right) = B.$$

We thus have $A_\theta(t_\theta + \delta) < B < A_\theta(t_\theta)$. Because $A_\theta(\cdot)$ is concave, this implies that every maximizer of A_θ has to be strictly smaller than $t_\theta + \delta$ which proves (5.4). From (5.4), we immediately have

$$\begin{aligned} & \mathbb{P}_\theta\left\{\|\widehat{\theta} - \theta\|_2 \geq t_\theta + \delta\right\} \\ \leq \; & \mathbb{P}_\theta\left\{\mathbb{E}_\theta A_\theta(t_\theta) - A_\theta(t_\theta) \geq \frac{\Delta}{2}\right\} + \mathbb{P}_\theta\left\{A_\theta(t_\theta + \delta) - \mathbb{E}_\theta A_\theta(t_\theta + \delta) \geq \frac{\Delta}{2}\right\} \end{aligned} \tag{5.6}$$

where

$$\Delta := \mathbb{E}_\theta A_\theta(t_\theta) - \mathbb{E}_\theta A_\theta(t_\theta + \delta).$$

We now note that $\mathbb{E}_\theta A_\theta(t) = m_\theta(t) - t^2/2 = G_\theta(t)$ where $G_\theta(t)$ and $m_\theta(t)$ are defined in (2.1). Therefore, from Lemma 2.2(4), we get

$$\Delta = G_\theta(t_\theta) - G_\theta(t_\theta + \delta) \geq \frac{\delta^2}{2}.$$

Thus (5.6) gives

$$\begin{aligned} & \mathbb{P}_\theta\left\{\|\widehat{\theta} - \theta\|_2 \geq t_\theta + \delta\right\} \\ \leq \; & \mathbb{P}_\theta\left\{\mathbb{E}_\theta A_\theta(t_\theta) - A_\theta(t_\theta) \geq \frac{\delta^2}{4}\right\} + \mathbb{P}\left\{A_\theta(t_\theta + \delta) - \mathbb{E}_\theta A_\theta(t_\theta + \delta) \geq \frac{\delta^2}{4}\right\} \end{aligned}$$

We now use the trivial fact that for every $t \geq 0$, the quantity $A_\theta(t)$, as a function of Z, is Lipschitz with Lipschitz constant t. Therefore, by standard concentration for Lipschitz functions of Gaussian random variables, we get

$$\mathbb{P}_\theta \left\{ \|\widehat{\theta} - \theta\|_2 \geq t_\theta + \delta \right\} \leq \exp\left(\frac{-\delta^4}{32 t_\theta^2}\right) + \exp\left(\frac{-\delta^4}{32(t_\theta+\delta)^2}\right) \leq 2\exp\left(\frac{-\delta^4}{32(t_\theta+\delta)^2}\right)$$

which proves inequality (2.2).

We now turn to the proof of (2.3). For convenience, let $L := \|\widehat{\theta}(X;\Theta,f) - \theta\|_2$. First assume that $t_\theta \geq 1$. Using (2.2), we can write

$$\mathbb{P}_\theta \left\{ L \geq t_\theta + x\sqrt{t_\theta} \right\} \leq 2\exp\left(\frac{-x^4}{32\left(1 + x t_\theta^{-1/2}\right)^2}\right) \leq 2\exp\left(\frac{-x^4}{32(1+x)^2}\right).$$

As a result, via the identity $\mathbb{E}X_+^2 = 2\int_0^\infty x\mathbb{P}\{X \geq x\}dx$ which holds for every random variable X, we obtain

$$\begin{aligned}\mathbb{E}_\theta \left(\frac{L - t_\theta}{\sqrt{t_\theta}}\right)_+^2 &= 2\int_0^\infty x\ \mathbb{P}_\theta \left\{ L \geq t_\theta + x\sqrt{t_\theta} \right\} dx \\ &\leq 4\int_0^\infty x\ \exp\left(\frac{-x^4}{32(1+x)^2}\right) dx \leq 84.\end{aligned}$$

where we have also used the fact that the integral above is at most 21 (as can be verified by numerical computation). Note that the above bound also implies that

$$\mathbb{E}_\theta \left(\frac{L - t_\theta}{\sqrt{t_\theta}}\right)_+ \leq \sqrt{84}.$$

Thus if $L := \|\widehat{\theta}(X;\Theta,f) - \theta\|_2$, then the inequality

$$L \leq (L - t_\theta)_+^2 + t_\theta^2 + 2t_\theta(L - t_\theta)_+ \tag{5.7}$$

together with the above two bounds for $\mathbb{E}(L - t_\theta)_+^2$ and $\mathbb{E}_\theta(L - t_\theta)_+$ proves (2.3) in the case when $t_\theta \geq 1$.

When $t_\theta \leq 1$, inequality (2.2) gives

$$\mathbb{P}_\theta \left\{ L \geq t_\theta + x \right\} \leq 2\exp\left(\frac{-x^4}{32(t+x)^2}\right) \leq 2\exp\left(\frac{-x^4}{32(1+x)^2}\right)$$

which implies, as before, that

$$\mathbb{E}_\theta (L - t_\theta)_+^2 \leq 84 \quad \text{and} \quad \mathbb{E}_\theta (L - t_\theta)_+ \leq \sqrt{84}.$$

so that the required inequality (2.3) again follows from (5.7). □

Acknowledgment

We are thankful to Sivaraman Balakrishnan for helpful discussions.

References

Bühlmann, P. and S. van De Geer (2011). *Statistics for high-dimensional data: methods, theory and applications*. Springer Science & Business Media.

Chatterjee, S. (2014). A new perspective on least squares under convex constraint. *The Annals of Statistics 42*(6), 2340–2381.

Chen, X., A. Guntuboyina, and Y. Zhang (2016). On Bayes risk lower bounds. *Journal of Machine Learning Research 17*, 1–58.

Groeneboom, P. and G. Jongbloed (2014). *Nonparametric Estimation under Shape Constraints: Estimators, Algorithms and Asymptotics*, Volume 38. Cambridge University Press.

Le Cam, L. (1973). Convergence of estimates under dimensionality restrictions. *Annals of Statistics 1*, 38–53.

Lehmann, E. L. and G. Casella (1998). *Theory of Point Estimation* (2nd ed.). New York: Springer.

Muro, A. and S. van de Geer (2015). Concentration behavior of the penalized least squares estimator. *arXiv preprint arXiv:1511.08698*.

Tsybakov, A. (2009). *Introduction to Nonparametric Estimation*. Springer-Verlag.

van de Geer, S. and M. Wainwright (2015). On concentration for (regularized) empirical risk minimization. *arXiv preprint arXiv:1512.00677*.

Woodroofe, M. and J. Sun (1993). A penalized maximum likelihood estimate of f (0+) when f is non-increasing. *Statistica Sinica*, 501–515.

Zhang, L. (2013). Nearly optimal minimax estimator for high-dimensional sparse linear regression. *The Annals of Statistics 41*(4), 2149–2175.